

## Parallel Solutions Inc

A great deal of research is being done in the areas of artificial vision and neural networks. Although much of this research has been theoretical in nature, many of the techniques developed through these efforts are now mature enough for use in practical applications.

Automated Visual Inspection Using Artificial Neural Networks explains the application of recently emerging technology in the areas of artificial vision and neural networks to automated visual inspection. The information is organised in a clear, informative manner, bridging the gap between theoretical research and practical application. Significantly this book includes: \* broad coverage of all aspects of the automated visual inspection problem, \* details of the HAVENET neural network and the CAMERA vision model, and \* detailed descriptions of practical applications of intelligent visual inspection.

Plastic films are high-performance materials which play an essential part in modern life. The plastics films industry uses state-of-the-art manufacturing processes and is continuously seeking out new technologies to improve its performance. The understanding of the nature of plastic films, their production techniques, applications and their characterisation is essential for producing new types of plastic films. This handbook has been written to discuss the production and main uses of plastic films. Plastic films are high-performance materials which play an essential part in modern life. Plastic films are mostly used in packaging applications but as will be

seen from this book they are also used in the agricultural, medical and engineering fields. The plastics films industry uses state-of-the-art manufacturing processes and is continuously seeking out new technologies to improve its performance. The understanding of the nature of plastic films, their production techniques, applications and their characterisation is essential for producing new types of plastic films. This handbook has been written to discuss the production and main uses of plastic films.

Computer architecture deals with the physical configuration, logical structure, formats, protocols, and operational sequences for processing data, controlling the configuration, and controlling the operations over a computer. It also encompasses word lengths, instruction codes, and the interrelationships among the main parts of a computer or group of computers. This two-volume set offers a comprehensive coverage of the field of computer organization and architecture.

For more than 40 years, Computerworld has been the leading source of technology news and information for IT influencers worldwide. Computerworld's award-winning Web site ([Computerworld.com](http://Computerworld.com)), twice-monthly publication, focused conference series and custom research form the hub of the world's largest global IT media network.

An Introduction to Parallel Programming, Second Edition presents a tried-and-true tutorial approach that shows students how to develop effective parallel programs with MPI, Pthreads and OpenMP. As the first undergraduate text to directly address compiling and running parallel

programs on multi-core and cluster architecture, this second edition carries forward its clear explanations for designing, debugging and evaluating the performance of distributed and shared-memory programs while adding coverage of accelerators via new content on GPU programming and heterogeneous programming. New and improved user-friendly exercises teach students how to compile, run and modify example programs. Takes a tutorial approach, starting with small programming examples and building progressively to more challenging examples Explains how to develop parallel programs using MPI, Pthreads and OpenMP programming models A robust package of online ancillaries for instructors and students includes lecture slides, solutions manual, downloadable source code, and an image bank New to this edition: New chapters on GPU programming and heterogeneous programming New examples and exercises related to parallel algorithms

"An overview of the multidisciplinary field of data mining, this book focuses specifically on new methodologies and case studies. Included are case studies written by 44 leading scientists and talented young scholars from seven different countries. Topics covered include data mining based on rough sets, the impact of missing data, and mining free text for structure. In addition, the four basic mining operations supported by numerous mining techniques are addressed: predictive model creation supported by supervised induction techniques; link analysis supported by association discovery and sequence discovery techniques; DB segmentation supported by clustering techniques; and deviation

detection supported by statistical techniques." Massively Parallel Systems (MPSs) with their scalable computation and storage space promises are becoming increasingly important for high-performance computing. The growing acceptance of MPSs in academia is clearly apparent. However, in industrial companies, their usage remains low. The programming of MPSs is still the big obstacle, and solving this software problem is sometimes referred to as one of the most challenging tasks of the 1990's. The 1994 working conference on "Programming Environments for Massively Parallel Systems" was the latest event of the working group WG 10.3 of the International Federation for Information Processing (IFIP) in this field. It succeeded the 1992 conference in Edinburgh on "Programming Environments for Parallel Computing". The research and development work discussed at the conference addresses the entire spectrum of software problems including virtual machines which are less cumbersome to program; more convenient programming models; advanced programming languages, and especially more sophisticated programming tools; but also algorithms and applications.

Company Profiles: Parallel Solutions Inc  
Parallel Solution of Partial Differential Equations  
Springer Science & Business Media

Foreword by Bjarne Stroustrup Software is generally

acknowledged to be the single greatest obstacle preventing mainstream adoption of massively-parallel computing. While sequential applications are routinely ported to platforms ranging from PCs to mainframes, most parallel programs only ever run on one type of machine. One reason for this is that most parallel programming systems have failed to insulate their users from the architectures of the machines on which they have run. Those that have been platform-independent have usually also had poor performance. Many researchers now believe that object-oriented languages may offer a solution. By hiding the architecture-specific constructs required for high performance inside platform-independent abstractions, parallel object-oriented programming systems may be able to combine the speed of massively-parallel computing with the comfort of sequential programming. Parallel Programming Using C++ describes fifteen parallel programming systems based on C++, the most popular object-oriented language of today. These systems cover the whole spectrum of parallel programming paradigms, from data parallelism through dataflow and distributed shared memory to message-passing control parallelism. For the parallel programming community, a common parallel application is discussed in each chapter, as part of the description of the system itself. By comparing the implementations of the polygon overlay problem in

each system, the reader can get a better sense of their expressiveness and functionality for a common problem. For the systems community, the chapters contain a discussion of the implementation of the various compilers and runtime systems. In addition to discussing the performance of polygon overlay, several of the contributors also discuss the performance of other, more substantial, applications. For the research community, the contributors discuss the motivations for and philosophy of their systems. As well, many of the chapters include critiques that complete the research arc by pointing out possible future research directions. Finally, for the object-oriented community, there are many examples of how encapsulation, inheritance, and polymorphism can be used to control the complexity of developing, debugging, and tuning parallel software.

Brings together, analyzes, and contextualizes the latest findings and practical applications

Polyphosphazenes, an emerging class of polymers, include macromolecules, which have been proven to be biocompatible, biodegradable, and bioactive.

Their unprecedented structural diversity and unique properties make them suitable as vaccine adjuvants, microencapsulating agents, biodegradable materials, scaffolds for tissue engineering, biocompatible coatings, and carriers for gene delivery.

Polyphosphazenes for Biomedical Applications

offers a thorough review of polyphosphazene research findings in the life sciences, chemistry, and chemical engineering. It emphasizes biomedical applications as well as recent advances in polyphosphazene development such as high-throughput discovery and the latest controlled methods of synthesis. The book brings together, analyzes, and contextualizes a wealth of knowledge that previously could only be found scattered throughout the scientific literature. Following two introductory chapters, the book reviews: Vaccine delivery and immunomodulation Biomaterials Drug delivery systems Biodetection Well-defined polyphosphazenes: synthetic aspects and novel molecular architectures All the chapters have been written by leading researchers in the field. Editor Alexander Andrianov, who has led the effort to commercialize polyphosphazenes for biomedical applications, has carefully reviewed and edited all chapters to ensure readability, accuracy, and thoroughness. Polyphosphazenes for Biomedical Applications is not only intended for researchers working in polyphosphazene chemistry, but also for all researchers seeking solutions to problems arising in the areas of biomaterials, drug delivery systems, and controlled release formulations.

Asia has become the new battle ground for the war against HIV/AIDS. The magnitude of the potential public health problems caused by AIDS in this

populous continent may become a catastrophic disaster. A 10% rate of prevalence of HIV-1 in India and China alone would mean more than 200 million people are infected with HIV. AIDS in Asia is useful as a comprehensive, up-to-date AIDS reference book for public health and medical professionals. This volume provides concrete information on the diagnosis, treatment, care, prevention and impact of AIDS. Part I contains 'Snapshots of HIV/AIDS in Asia.' Countries and regions included in this section are: Thailand, India, China, Japan, Vietnam, Laos, Cambodia, Myanmar, Malaysia, Indonesia, the Philippines, Pakistan, Nepal, Sri Lanka, Hawaii and the Pacific Islands, Hong Kong and Taiwan. Part II addresses the molecular epidemiology of HIV/AIDS. Part III deals with the diagnosis, surveillance and projected scenarios of the AIDS epidemic. Part IV outlines prevention efforts and treatment options. Part V provides an overview of the ongoing collaborative efforts involved in several different nations in the worldwide war against AIDS. This volume will be invaluable to all the public health professionals and researchers working in this field. "...the book is a useful addition to the HIV/AIDS literature." "AIDS in Asia offers a comprehensive, interesting overview of the epidemic there and of general issues that will influence its progression." -Roger Detels, MD, MS, University of California-Los Angeles The Journal of the American Medical

Association, Book Review, 293:15

Discover how to streamline complex bioinformatics applications with parallel computing This publication enables readers to handle more complex bioinformatics applications and larger and richer data sets. As the editor clearly shows, using powerful parallel computing tools can lead to significant breakthroughs in deciphering genomes, understanding genetic disease, designing customized drug therapies, and understanding evolution. A broad range of bioinformatics applications is covered with demonstrations on how each one can be parallelized to improve performance and gain faster rates of computation. Current parallel computing techniques and technologies are examined, including distributed computing and grid computing. Readers are provided with a mixture of algorithms, experiments, and simulations that provide not only qualitative but also quantitative insights into the dynamic field of bioinformatics. Parallel Computing for Bioinformatics and Computational Biology is a contributed work that serves as a repository of case studies, collectively demonstrating how parallel computing streamlines difficult problems in bioinformatics and produces better results. Each of the chapters is authored by an established expert in the field and carefully edited to ensure a consistent approach and high standard throughout the publication. The work is

organized into five parts: \* Algorithms and models \* Sequence analysis and microarrays \* Phylogenetics \* Protein folding \* Platforms and enabling technologies Researchers, educators, and students in the field of bioinformatics will discover how high-performance computing can enable them to handle more complex data sets, gain deeper insights, and make new discoveries.

Content Description #Includes bibliographical references and index.

This IMA Volume in Mathematics and its Applications ALGORITHMS FOR PARALLEL PROCESSING is based on the proceedings of a workshop that was an integral part of the 1996-97 IMA program on "MATHEMATICS IN HIGH-PERFORMANCE COMPUTING. " The workshop brought together algorithm developers from theory, combinatorics, and scientific computing. The topics ranged over models, linear algebra, sorting, randomization, and graph algorithms and their analysis. We thank Michael T. Heath of University of Illinois at Urbana (Computer Science), Abhiram Ranade of the Indian Institute of Technology (Computer Science and Engineering), and Robert S. Schreiber of Hewlett Packard Laboratories for their excellent work in organizing the workshop and editing the proceedings. We also take this opportunity to thank the National Science Foundation (NSF) and the Army Research Office (ARO), whose financial support made the workshop possible. A vner Friedman Robert Gulliver v PREFACE The Workshop on Algorithms for Parallel Processing was held at the IMA September 16 - 20, 1996; it was the first workshop of the IMA year dedicated to the mathematics of high performance computing. The workshop organizers were Abhiram Ranade of The Indian Institute

of Technology, Bombay, Michael Heath of the University of Illinois, and Robert Schreiber of Hewlett Packard Laboratories. Our idea was to bring together researchers who do innovative, exciting, parallel algorithms research on a wide range of topics, and by sharing insights, problems, tools, and methods to learn something of value from one another.

Contributed presentations were given by over 50 researchers representing the state of parallel CFD art and architecture from Asia, Europe, and North America. Major developments at the 1999 meeting were: (1) the effective use of as many as 2048 processors in implicit computations in CFD, (2) the acceptance that parallelism is now the 'easy part' of large-scale CFD compared to the difficulty of getting good per-node performance on the latest fast-clocked commodity processors with cache-based memory systems, (3) favorable prospects for Lattice-Boltzmann computations in CFD (especially for problems that Eulerian and even Lagrangian techniques do not handle well, such as two-phase flows and flows with exceedingly multiple-connected domains with a lot of holes in them, but even for conventional flows already handled well with the continuum-based approaches of PDEs), and (4) the nascent integration of optimization and very large-scale CFD. Further details of Parallel CFD'99, as well as other conferences in this series, are available at

<http://www.parcfd.org>

This IMA Volume in Mathematics and its Applications PARALLEL SOLUTION OF PARTIAL DIFFERENTIAL EQUATIONS is based on the proceedings of a workshop with the same title. The workshop was an integral part of the 1996-97 IMA program on "MATHEMATICS IN HIGH-PERFORMANCE COMPUTING." I would like to thank Petter Bjørstad of the Institutt for Informatikk, University of Bergen and Mitchell Luskin of the School of Mathematics, University of Minnesota for their excellent work as organizers of the

meeting and for editing the proceedings. I also take this opportunity to thank the National Science Foundation (NSF), Department of Energy (DOE), and the Army Research Office (ARO), whose financial support made the workshop possible. Willard Miller, Jr., Professor and Director

### PREFACE

The numerical solution of partial differential equations has been of major importance to the development of many technologies and has been the target of much of the development of parallel computer hardware and software. Parallel computers offer the promise of greatly increased performance and the routine calculation of previously intractable problems. The papers in this volume were presented at the IMA workshop on the Parallel Solution of PDE held during June 9-13, 1997. The workshop brought together leading numerical analysts, computer scientists, and engineers to assess the state-of-the-art and to consider future directions.

This book constitutes the thoroughly refereed post-conference proceedings of the 8th International Conference on High Performance Computing for Computational Science, VECPAR 2008, held in Toulouse, France, in June 2008. The 51 revised full papers presented together with the abstract of a surveying and look-ahead talk were carefully reviewed and selected from 73 submissions. The papers are organized in topical sections on parallel and distributed computing, cluster and grid computing, problem solving environment and data centric, numerical methods, linear algebra, computing in geosciences and biosciences, imaging and graphics, computing for aerospace and engineering, and high-performance data management in grid environments. The use of parallel programming and architectures is essential for simulating and solving problems in modern computational practice. There has been rapid progress in microprocessor architecture, interconnection technology and software development, which are influencing directly the rapid

growth of parallel and distributed computing. However, in order to make these benefits usable in practice, this development must be accompanied by progress in the design, analysis and application aspects of parallel algorithms. In particular, new approaches from parallel numerics are important for solving complex computational problems on parallel and/or distributed systems. The contributions to this book are focused on topics most concerned in the trends of today's parallel computing. These range from parallel algorithmics, programming, tools, network computing to future parallel computing. Particular attention is paid to parallel numerics: linear algebra, differential equations, numerical integration, number theory and their applications in computer simulations, which together form the kernel of the monograph. We expect that the book will be of interest to scientists working on parallel computing, doctoral students, teachers, engineers and mathematicians dealing with numerical applications and computer simulations of natural phenomena.

"I could not stop reading and highlighting! Thank you for writing it! Finally someone that understands and gets that it is impossible to coparent with a high conflict person!" -L. Deshea

"This information was liberting to hear...this is exactly what I am and have been going through. I have purchased your book on Amazon and started reading it. I want to offer a sincere thank you...I am learning a lot about aligning myself with the parallel parenting ideal. Thank you for your posts and for your book. It has truly changed my life." -S. Delgado

"Thank you for writing this book. It is good to not only know that I am not alone in this, but to have a safe, healthy plan for my son and myself moving forward." -J. Dillard

"I want to see these changes for all families in my lifetime!" -A. Ngyen

The Parallel Parenting Solution was written for people who want to eliminate the unnecessary drama of Coparenting with a high-conflict ex-particularly in the wake of a hard-fought divorce

battle. Coparenting does not have to take a toll on mental health, deplete your finances, compromise your values, or threaten your safety. For those experiencing the effects of Trendy-Trendy Coparenting as a living hell, your experience is valid. We've been there and lived it. But have no fear. There is hope, and it's called Parallel Parenting. It's based on the premise that all parties can achieve the highest outcomes for themselves and their families when they are free to work in parallel, rather than being thrown into the chaotic emotional enmeshment soup that is Trendy-Trendy Coparenting. Understanding the exploitative and conflict-producing fantasy known as Trendy-Trendy Coparenting as sold to us by the divorce industry vultures is as important as understanding the down-to-earth tactics of how to deal with your high-conflict ex. This book will teach you both. Families who want to heal after a divorce should not start with pie-in-the-sky expectations. They should start in reality. That's where Parallel Parenting exists, and we hope you will join us there.

This book constitutes the refereed proceedings of the Third International Euro-Par Conference, held in Passau, Germany, in August 1997. The 178 revised papers presented were selected from more than 300 submissions on the basis of 1101 reviews. The papers are organized in accordance with the conference workshop structure in tracks on support tools and environments, routing and communication, automatic parallelization, parallel and distributed algorithms, programming languages, programming models and methods, numerical algorithms, parallel architectures, HPC applications, scheduling and load balancing, performance evaluation, instruction-level parallelism, database systems, symbolic computation, real-time systems, and an ESPRIT workshop.

Although the origins of parallel computing go back to the last century, it was only in the 1970s that parallel and vector

computers became available to the scientific community. The first of these machines-the 64 processor Illiac IV and the vector computers built by Texas Instruments, Control Data Corporation, and then CRA Y Research Corporation-had a somewhat limited impact. They were few in number and available mostly to workers in a few government laboratories. By now, however, the trickle has become a flood. There are over 200 large-scale vector computers now installed, not only in government laboratories but also in universities and in an increasing diversity of industries. Moreover, the National Science Foundation's Super computing Centers have made large vector computers widely available to the academic community. In addition, smaller, very cost-effective vector computers are being manufactured by a number of companies. Parallelism in computers has also progressed rapidly. The largest super computers now consist of several vector processors working in parallel. Although the number of processors in such machines is still relatively small (up to 8), it is expected that an increasing number of processors will be added in the near future (to a total of 16 or 32). Moreover, there are a myriad of research projects to build machines with hundreds, thousands, or even more processors. Indeed, several companies are now selling parallel machines, some with as many as hundreds, or even tens of thousands, of processors.

This book constitutes the refereed proceedings of the 7th International Workshop on Advanced Parallel Processing Technologies, APPT 2007, held in Guangzhou, China, in November 2007. The 78 revised full papers presented were carefully reviewed and selected from 346 submissions. All current aspects in parallel and distributed computing are addressed ranging from hardware and software issues to algorithmic aspects and advanced applications. The papers are organized in topical sections.

This volume reviews, in the context of partial differential equations, algorithm development that has been specifically aimed at computers that exhibit some form of parallelism. Emphasis is on the solution of PDEs because these are typically the problems that generate high computational demands. The authors discuss architectural features of these computers inasmuch as they influence algorithm performance, and provide insight into algorithm characteristics that allow effective use of hardware.

InfoWorld is targeted to Senior IT professionals. Content is segmented into Channels and Topic Centers. InfoWorld also celebrates people, companies, and projects.

This monograph evolved from my Ph. D dissertation completed at the Laboratory of Computer Science, MIT, during the Summer of 1986. In my dissertation I proposed a pipelined code mapping scheme for array operations on static dataflow architectures. The main addition to this work is found in Chapter 12, reflecting new research results developed during the last three years since I joined McGill University—results based upon the principles in my dissertation. The terminology dataflow software pipelining has been consistently used since publication of our 1988 paper on the argument-fetching dataflow architecture model at McGill University [43]. In the first part of this book we describe the static data flow graph model as an operational model for concurrent computation. We look at timing considerations for program graph execution on an ideal static dataflow computer, examine the notion of pipe lining, and characterize its performance. We discuss balancing techniques used to transform certain graphs into fully pipelined data flow graphs. In particular, we show how optimal balancing of an acyclic data flow graph can be formulated as a linear programming problem for which an optimal solution exists. As a major result, we show the optimal balancing problem of acyclic data

flow graphs is reduceable to a class of linear programming problem, the net work flow problem, for which well-known efficient algorithms exist. This result disproves the conjecture that such problems are computationally hard.

Parallel-Algorithms for Regular Architectures is the first book to concentrate exclusively on algorithms and paradigms for programming parallel computers such as the hypercube, mesh, pyramid, and mesh-of-trees. Algorithms are given to solve fundamental tasks such as sorting and matrix operations, as well as problems in the field of image processing, graph theory, and computational geometry. The first chapter defines the computer models, problems to be solved, and notation that will be used throughout the book. It also describes fundamental abstract data movement operations that serve as the foundation to many of the algorithms presented in the book. The remaining chapters describe efficient implementations of these operations for specific models of computation and present algorithms (with asymptotic analyses) that are often based on these operations. The algorithms presented are the most efficient known, including a number of new algorithms for the hypercube and mesh-of-trees that are better than those that have previously appeared in the literature. The chapters may be read independently, allowing anyone interested in a specific model to read the introduction and then move directly to the chapter(s) devoted to the particular model of interest. Russ Miller is Assistant Professor in the Department of Computer Science, State University of New York at Buffalo. Quentin F. Stout is Associate Professor in the Department of Electrical Engineering and Computer Science at the University of Michigan. Parallel Algorithms for Regular Architectures is included in the Scientific Computation series, edited by Dennis Gannon.

Annotation. This book constitutes the refereed proceedings of

the 11th International Conference on High-Performance Computing, HiPC 2004, held in Bangalore, India in December 2004. The 48 revised full papers presented were carefully reviewed and selected from 253 submissions. The papers are organized in topical sections on wireless network management, compilers and runtime systems, high performance scientific applications, peer-to-peer and storage systems, high performance processors and routers, grids and storage systems, energy-aware and high-performance networking, and distributed algorithms.

A comprehensive overview of OpenMP, the standard application programming interface for shared memory parallel computing—a reference for students and professionals. "I hope that readers will learn to use the full expressibility and power of OpenMP. This book should provide an excellent introduction to beginners, and the performance section should help those with some experience who want to push OpenMP to its limits." —from the foreword by David J. Kuck, Intel Fellow, Software and Solutions Group, and Director, Parallel and Distributed Solutions, Intel Corporation

OpenMP, a portable programming interface for shared memory parallel computers, was adopted as an informal standard in 1997 by computer scientists who wanted a unified model on which to base programs for shared memory systems. OpenMP is now used by many software developers; it offers significant advantages over both hand-threading and MPI. Using OpenMP offers a comprehensive introduction to parallel programming concepts and a detailed overview of OpenMP. Using OpenMP discusses hardware developments, describes where OpenMP is applicable, and compares OpenMP to other programming interfaces for shared and distributed memory parallel architectures. It introduces the individual features of OpenMP, provides many source code examples that demonstrate the use and functionality of the language

constructs, and offers tips on writing an efficient OpenMP program. It describes how to use OpenMP in full-scale applications to achieve high performance on large-scale architectures, discussing several case studies in detail, and offers in-depth troubleshooting advice. It explains how OpenMP is translated into explicitly multithreaded code, providing a valuable behind-the-scenes account of OpenMP program performance. Finally, Using OpenMP considers trends likely to influence OpenMP development, offering a glimpse of the possibilities of a future OpenMP 3.0 from the vantage point of the current OpenMP 2.5. With multicore computer use increasing, the need for a comprehensive introduction and overview of the standard interface is clear. Using OpenMP provides an essential reference not only for students at both undergraduate and graduate levels but also for professionals who intend to parallelize existing codes or develop new parallel programs for shared memory computer architectures.

PCMag.com is a leading authority on technology, delivering Labs-based, independent reviews of the latest products and services. Our expert industry analysis and practical solutions help you make better buying decisions and get more from technology.

[Copyright: ad8e07f1b7675e156cd4b7c9138e610d](https://www.pcmag.com/)